

Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 9, 2012

B. Niven-Jenkins, Ed.
Velocix (Alcatel-Lucent)
G. Watson
BT
N. Bitar
Verizon
J. Medved
Juniper Networks
S. Previdi
Cisco Systems
December 7, 2011

Use Cases for ALTO within CDNs
draft-jenkins-alto-cdn-use-cases-02

Abstract

For some time, Content Distribution Networks (CDNs) have been used in the delivery of some Internet services (e.g. delivery of websites, software updates and video delivery) as they provide numerous benefits including reduced delivery cost for cacheable content, improved quality of experience for end users and increased robustness of delivery.

In order to derive the optimal benefit from a CDN it is preferable to deliver content from the servers (caches) that are "closest" to the End User requesting the content, where "closest" may be as simple as "geographical or network distance" combined with CDN server load within a location, but may also consider other more complex combinations of metrics and CDN or Network Service Provider (NSP) policies.

There are a number of different ways in which a CDN may obtain the necessary network topology and/or cost information to allow it to serve End Users from the most optimal servers/locations, such as static configuration, passively listening to routing protocols directly, active probing of underlying network(s), or obtaining topology and cost by querying an information service such as the ALTO map & cost services.

This document describes the use cases for a CDN to be able to obtain network topology and cost information from an ALTO server(s).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 9, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	4
1.1.	Terminology	5
2.	CDN overview	5
3.	CDN & ALTO Use Cases	7
3.1.	Exposing NSP End User Reachability to a CDN	8
3.2.	Exposing CDN End User Reachability to CSPs	9
3.3.	CDN deployed within a Broadband network	10
3.4.	CDN delivering Over-The-Top of a NSP's network	11
3.5.	CDN acquiring content from multiple upstream sources (Origins)	11
3.6.	Additional Use Cases	12
4.	IANA Considerations	13
5.	Security Considerations	13
6.	Contributing Authors	13
7.	Acknowledgements	13
8.	Normative References	13
	Authors' Addresses	14

1. Introduction

For some time, Content Distribution Networks (CDNs) have been used in the delivery of some Internet services (e.g. delivery of websites, software updates and video delivery) as they provide numerous benefits including reduced delivery cost for cacheable content, improved quality of experience for end users and increased robustness of delivery.

A CDN typically consists of a network of servers often attached to Network Service Provider (NSP) networks. The point of attachment is often as close to content consumers and peering points as economically or operationally feasible in order to decrease traffic load on the NSP backbone and to provide better user experience measured by reduced latency and higher throughput.

As the volume of video and multimedia content delivered over the Internet is rapidly increasing and expected to continue doing so in the future, existing CDN providers are scaling up their infrastructure and many NSPs are deploying their own CDNs. The result of such deployments is typically that more CDN servers are being deployed within NSP networks and those CDN servers are being deployed in locations that are "deeper" (i.e. geographically closer to the NSP's End Users) than was previously the case.

In order to derive the optimal benefit from a CDN it is preferable to deliver content from the servers (caches) that are "closest" to the End User requesting the content, where "closest" may be as simple as "geographical or network distance" combined with CDN server load within a location, but may also consider other more complex combinations of metrics and CDN or NSP policies.

When CDN servers are deployed outside of an NSP's network or in a small number of central locations within an NSP's network a simplified view of the NSP's topology or an approximation of proximity is typically sufficient to enable the CDN to serve End Users from the optimal server/location. As CDN servers are deployed deeper within NSP networks it becomes necessary for the CDN to have more detailed knowledge of the underlying network topology and costs between network locations in order to enable the CDN to serve End Users from the most optimal servers for the NSP.

There are a number of different ways in which a CDN may obtain the necessary network topology and/or cost information to allow it to serve End Users from the most optimal servers/locations, such as static configuration, passively listening to routing protocols directly, active probing of underlying network(s), or obtaining topology and cost by querying an information service such as the ALTO

map & cost services.

The rest of this document describes the use cases for a CDN to be able to obtain network topology and cost information from an ALTO server(s).

1.1. Terminology

The following terms are taken from [[I-D.ietf-cdni-problem-statement](#)] and repeated here for completeness.

Content Distribution Network (CDN) / Content Delivery Network (CDN): Network infrastructure in which the network elements cooperate at layers 4 through layer 7 for more effective delivery of Content to User Agents. Typically a CDN consists of a Request Routing system, a Distribution System (that includes a set of Surrogates), a Logging System and a CDN control system.

Content Service Provider (CSP): Provides a Content Service to End Users (which the End Users access via a User Agent). A CSP may own the Content made available as part of the Content Service, or may license content rights from another party.

End User (EU): The 'real' user of the system, typically a human but maybe some combination of hardware and/or software emulating a human (e.g. for automated quality monitoring etc.)

Network Service Provider (NSP): Provides network-based connectivity/ services to Users.

Surrogate: A device/function that interacts with other elements of the CDN for the control and distribution of Content within the CDN and interacts with User Agents for the delivery of the Content.

User Agent (UA): Software (or a combination of hardware and software) through which the End User interacts with the Content Service. The User Agent will communicate with the CSP's Service for the selection of content and one or more CDNs for the delivery of the Content. Such communication is not restricted to HTTP and may be via a variety of protocols. Examples of User Agents (non-exhaustive) are: Browsers, Set Top Boxes (STB), Dedicated content applications (e.g. media players), etc.

2. CDN overview

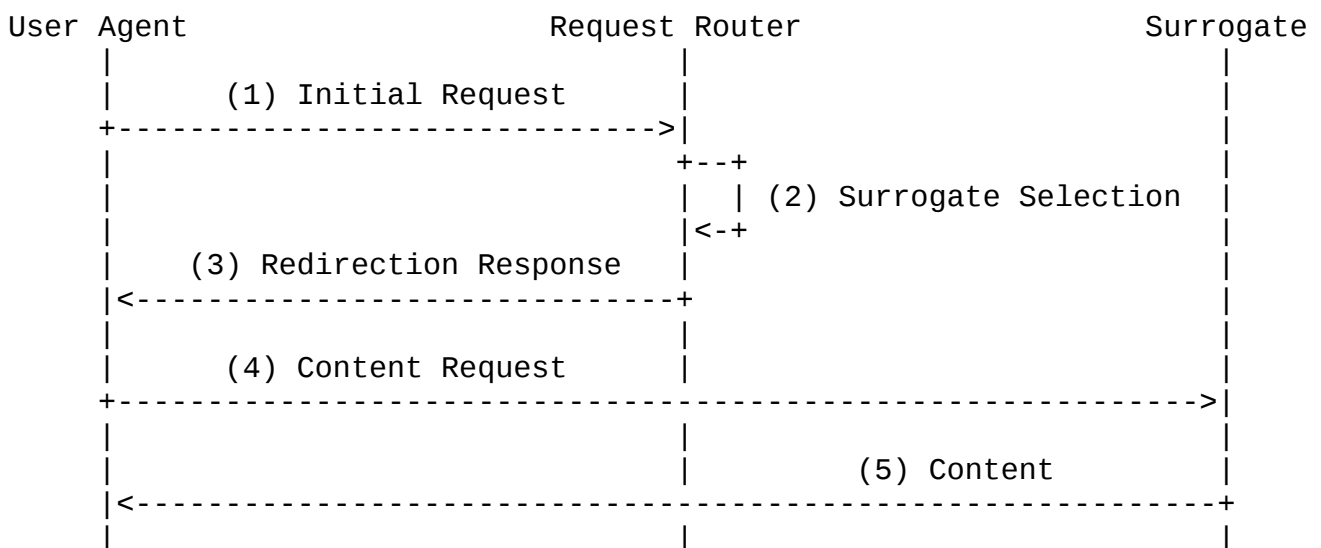
This section provides a high level and simplified overview of the operation of a CDN to help put the ALTO & CDN use cases into context.

A typical CDN consists of a number of functional components, however in the context of ALTO three functional components are of interest: The Request Routing function, the Surrogate (i.e. caching) function and the Origin function.

The Request Routing function within a CDN is responsible for receiving content requests from User Agents, obtaining and maintaining necessary information about a set of candidate Surrogates, and for selecting and redirecting the User Agent to the appropriate Surrogate.

The Surrogate function interacts with other elements of the CDN for the control and distribution of Content within the CDN and interacts with User Agents for the delivery of the Content.

The figure below shows a high level call flow showing the interaction between a User Agent, Request Router and Surrogate for the delivery of content in a single CDN.



1. The User Agent makes an initial request to the CDN. Depending on the type of content being delivered and the configuration of the CDN this request may be an application (e.g. HTTP, RTMP, etc.) level request directly from the User Agent or may be a DNS request via the User Agent's assigned DNS proxy.
2. The Request Router selects an appropriate Surrogate (or set of Surrogates) based on the User Agent's (or its proxy's) IP address, the Request Router's knowledge of the network topology and reachability cost between CDN caches and end users, and any additional CDN policies.

3. The Request Router responds to the UA's initial request with an appropriate response containing a redirection to the selected cache, for example by returning an appropriate DNS A/AAAA record, a HTTP 302 redirect, etc.
4. The User Agent uses the information provided in the Redirection Response to connect directly to the Surrogate and request the desired content.
5. If CDN policy allows the User Agent to receive the requested content, the Surrogate delivers the content to the User Agent.
 - A. [Not Shown] If the Surrogate does not have a copy of the requested content then it obtains it from the appropriate Origin Server.

Note: A Surrogate may not communicate with the Origin directly and instead obtain the requested content from other surrogates or caching layers in the CDN hierarchy. The details of how content requests filter through the CDN hierarchy to the Origin are internal to a specific CDN and are out of scope of this document.

3. CDN & ALTO Use Cases

The primary use case for ALTO in a CDN context is to improve the selection of a CDN Surrogate or Origin. The CDN makes use of an ALTO server to choose a better CDN Surrogate or Origin than would otherwise be the case. In its simplest form an ALTO server would provide an NSP with the capability to offer a service to a CDN which provides network map and cost information that the CDN can use to enhance its surrogate and/or Origin selection.

Although it is possible to obtain raw network map and cost information in other ways, for example passively listening to the NSP's routing protocols, the use of an ALTO service to expose that information may provide additional control to the NSP over how their network map/cost is exposed. Additionally it may enable the NSP to maintain a functional separation between their routing plane and network map computation functions. This may be attractive for a number of reasons, for example:

- o The ALTO service could provide a filtered view of the network and/or cost map that relates to CDN locations and their proximity to end users, for example to allow the NSP to control the level of topology detail they are willing to share with the CDN.
- o The ALTO service could apply additional policies to the network map and cost information to provide a CDN-specific view of the network map/cost, for example to allow the NSP to encourage the CDN to use network links that would not ordinarily be preferred by a Shortest Path First routing calculation.

- o The routing plane may be operated and controlled by a different operational entity (even within a single NSP) to the CDN and the ALTO service could provide a layer of separation because:
 - * The CDN is not able to passively listen to routing protocols.
 - * The NSP is not willing to allow the CDN to passively listen to routing protocols, e.g. because the NSP is concerned the CDN may inadvertently interfere with the routing plane or because the routing plane and the CDN are operated by different operational entities/groups (including different entities within the same NSP).

The use cases in this document are not necessarily specific as to the relationship between the commercial/operational entity that "owns" the ALTO service and the commercial/operational entity that "owns" the CDN service as it is assumed that such relationships will be deployment specific. Although the ownership of each service may affect the level of topology detail that the ALTO service will be permitted to expose, it is assumed that the general requirements a CDN places on the ALTO service should not change provided that the ALTO server is able to expose sufficient topology for the CDN to make appropriate surrogate and/or Origin selection decisions.

In general, the ALTO service is expected to be operated by an entity or entities that wish to optimize or otherwise influence request routing decisions. Some, non-exhaustive, examples of such entities are:

- o The entity that operates the CDN's underlying network (e.g. the "CDN deployed within a Broadband network" described in [Section 3.3](#)).
- o An NSP that wishes to optimize over-the-top content delivery from a CDN that is deployed outside of its network (e.g. the "CDN delivering Over-The-Top of a NSP" described in [Section 3.4](#)).
- o An NSP (that may or may not operate a CDN) or a CDN that wishes to advertise which End Users are reachable via its network/CDN (e.g. the exposing "End User reachability" use cases described in [Section 3.1](#) and [Section 3.2](#)).

The following sections outline some specific, non-exhaustive, example use cases, which are subsets of the primary use case outlined above but applied to specific usage examples to demonstrate how a CDN could make use of ALTO services.

[3.1](#). Exposing NSP End User Reachability to a CDN

In order for a Request Router to be able to make surrogate selection decisions, the Request Router needs to have information on which End User IP subnets are reachable via which networks or network locations. The granularity of location information required depends

on the specific deployment of the CDN relative to the End Users. For example, an Over-The-Top CDN whose surrogates are deployed only within the Internet "backbone" may only require knowledge of which End User IP subnets are reachable via which NSPs' networks, whereas a CDN deployed within a particular NSP's network requires a finer granularity of knowledge, i.e. which End User IP subnets are reachable via which regions within that NSP's network.

Such reachability information is often available via dynamic routing protocols, however it is likely that in a number of deployment scenarios that peering of the routing plane of the network with a CDN would be deemed unacceptable (e.g. where the CDN is operated by an entity other than the NSP(s) operating the underlying network).

Provided that some common mapping between ALTO PIDs and network locations (or entire networks) is known to both the NSP and the CDN, the network map services offered by ALTO could be used to expose which End User IP subnets are reachable via a particular network or network locations in order to export End User reachability to a Request Router to enable the NSP to expose End User reachability while also giving the NSP the ability to control the granularity of any End User reachability to network location mapping while also avoiding routing plane peering between the NSP and the CDN.

3.2. Exposing CDN End User Reachability to CSPs

This use case is similar to the use case described in [Section 3.1](#) however in this case it is the CDN that wishes to expose which End User IP subnets the CDN is capable of delivering services to.

In some deployments a particular CDN may not have reachability to (or may not wish to offer services to) every End User IP subnet reachable via the global Internet, for example because the CDN is only deployed within certain networks or geographic regions and the CDN is either unable (due to lack of reachability) or unwilling (due to cost or policy) to serve all End Users reachable via the global Internet.

The reachability offered by a particular CDN may not include all the End User IP subnets that a particular CSP requires in order to serve all of that CSP's customers and therefore if the CSP wishes to make use of the services offered by a CDN that can only serve a subset of their customers the CSP must have knowledge of which End User IP subnets a particular CDN is able to serve, so that they can select an appropriate CDN to use to deliver their service to particular subsets of their customers.

In such cases, the network map services offered by ALTO could be used to expose to a CSP which End User IP subnets are reachable via a

particular CDN. In the case where the CDN is operated by an NSP using ALTO in this way could also enable the NSP to separate the exposure of End User subnets reachable via their CDN from those reachable via their underlying network.

3.3. CDN deployed within a Broadband network

In this use case an NSP is providing Broadband services to its customers and has deployed a CDN within its Broadband network to alleviate the cost and/or improve the User Experience of content services for its Broadband customers.

The topology of Broadband access/backhaul networks is often much more constrained than metro/core networks. If CDN surrogates are deployed within the access/backhaul network, for a given set of End Users, the NSP is likely to want to utilise the surrogates deployed in the same access/backhaul region as those End Users in preference to surrogates deployed within the metro/core or within other access/backhaul regions.

It is common for Broadband subscribers to obtain their IP addresses dynamically and in many deployments the IP subnets allocated to a particular access/backhaul region can change relatively frequently. For example new IP subnets are added as the subscriber base grows, IP subnets are moved from one Broadband product in the NSP's portfolio to another as customers migrate in order to optimise the NSP's IP address utilisation, or they are simply moved as part of IP address management, etc.

Additionally, in certain cases, CDN surrogates deployed in a particular network region may become overloaded, leading to the CDN selecting alternative surrogates in a different region of the network for content delivery. If this occurs, an NSP may wish to influence such a decision, for example because the NSP would prefer a surrogate to be selected that is deployed in the the next best (cost-wise to the NSP) location.

In order to meet the NSP's objective of utilising their CDN to constrain access/backhaul costs and/or improve User Experience it is important that the CDN is able to select the most appropriate surrogate for a given set of End User IP subnets. Although the network topology is often reasonably static, in networks where the IP subnets allocated to a Broadband region are changing relatively frequently, static configuration of End User IP Subnets to CDN surrogates is possible but some NSPs may consider the operational burden of having to update such static configuration too high and would prefer the CDN to be able to dynamically obtain network map and cost information.

The NSP could make use of an ALTO service to expose a cost mapping/ranking between End User IP subnets (within that NSP's network) and CDN surrogate IP subnets/locations to meet its requirements while avoiding static configuration or direct integration of the CDN into its IP routing plane and to avoid the CDN being required to implement network layer routing computations.

3.4. CDN delivering Over-The-Top of a NSP's network

In this use case a CDN is deployed within one or more NSPs' networks but is delivering content "Over-The-Top" into another NSP's network (which we will call NSP Z) where the CDN is not deployed.

The CDN is unlikely to have direct visibility of NSP Z's network topology and may have a choice of entry points into NSP Z's network from which it could serve content to NSP Z's End Users. For example because NSP Z has direct peering links with the CDN in a number of locations or NSP Z has transit and/or peering relationships with several other NSPs where the CDN is deployed. NSP Z may wish to influence the locations from which the CDN serves content based on some factor(s) that it does not wish to expose directly or that might change over time. For example the available transit/peering capacity in different locations, the cost of connectivity to different locations, etc.

For example, a CSP is using NSP A's CDN and another NSP (NSP Z) has peering with NSP A in Los Angeles and New York. NSP Z would like to influence which peering location NSP A's CDN delivers content out of for NSP Z's end users by using their knowledge of the peering capacity they have deployed in LA & NY and the capacity they have between those peering locations and groups of end users without directly exposing their internal topology to NSP A.

An NSP could make use of an ALTO service to expose a cost mapping/ranking between End User IP subnets (within that NSP's network) and entry points into that NSP's network in order to try to influence the locations from which the CDN serves content into that NSP's network.

3.5. CDN acquiring content from multiple upstream sources (Origins)

Before a surrogate within a CDN is able to deliver content to an End User it must first have a copy of the content that the End User is requesting. Content may be obtained by surrogates in advance of it being requested (pre-positioned) by End Users or it may be obtained by surrogates dynamically in response to End User requests for the content (on-demand).

The ultimate source of the content (i.e. where the 'master' copy is

permanently stored) is typically referred to as the content's Origin (or Origin Server), however CDNs often employ an internal hierarchy of caching layers so that surrogates do not necessarily obtain content directly from the Origin. Such a hierarchy provides a number of benefits, for example reducing the number of requests for content received by the Origin (and therefore reducing the scaling requirements on the Origin), more efficient use of the underlying network as fewer copies of the content is required to traverse the same network links, etc.

For a particular CSP's content service multiple, possibly independently addressable, Origins may be used for resiliency and the Origin(s) may be deployed in a distributed manner across multiple geographic locations.

For the rest of this use case "upstream source" is used to mean either the Origin itself as well as other sources of the content, for example another caching layer within the CDN that has (or will obtain on demand) a copy of the content but is not the actual Origin.

Therefore, for a particular item of content, a surrogate may have a choice of upstream sources (both internal to the CDN and external Origins) from which it could obtain the content.

When presented with a choice of upstream sources, a surrogate may utilise some combination of policy and heuristics to decide which upstream sources (and in which order) it should attempt to use to obtain the content. A CDN may wish to utilise network topology & cost information as one of the inputs into such a content source selection process, for example to weight upstream sources that are topologically close to the surrogate that requires the content.

Additionally, where the CDN is deployed within one or more NSP networks, an NSP may want to try to influence the choice of upstream sources, for example the NSP may prefer the CDN to use content sources that are deployed within that NSP's network or within networks with which it has direct peering agreements with over other content sources.

An NSP (or a CSP) could provide an ALTO service which a CDN could use to obtain network topology and/or cost/ranking information to use as an input into surrogates' selection decisions for content sources.

3.6. Additional Use Cases

The following additional use case may be relevant to ALTO and will be described in more detail in a future version of this document:

- o Inter-provider CDN / CDN Interconnect.

4. IANA Considerations

This document makes no specific request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

5. Security Considerations

TBD.

6. Contributing Authors

Reinaldo Penno
Juniper Networks
Email: rpenno@juniper.net

Richard Alimi
Google
Email: ralimi@google.com

Richard Yang
Yale University
Email: ryr@yale.edu

7. Acknowledgements

The authors would like to thank Vijay Gurbani and Volker Hilt for their review comments and contributions to the text.

8. Normative References

- [I-D.ietf-cdni-problem-statement]
Niven-Jenkins, B., Faucheur, F., and N. Bitar, "Content Distribution Network Interconnection (CDNI) Problem Statement", [draft-ietf-cdni-problem-statement-01](#) (work in progress), October 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

Authors' Addresses

Ben Niven-Jenkins (editor)
Velocix (Alcatel-Lucent)
326 Cambridge Science Park
Milton Road, Cambridge CB4 0WG
UK

Email: ben@velocix.com

Grant Watson
BT

Email: grant.watson@bt.com

Nabil Bitar
Verizon
40 Sylvan Road
Waltham, MA 02145
USA

Email: nabil.bitar@verizon.com

Jan Medved
Juniper Networks

Email: jmedved@juniper.net

Stefano Previdi
Cisco Systems

Email: sprevidi@cisco.com